The development and validation of the Community Hyper-Spectral Infrared Microwave Earth Retrieval Algorithm (CHIMERA): the path forward.

Chris Barnet
and Antonia Gambacorta (collaborator on proposal)
NASA Sounder Science Team Meeting
Thursday, Oct. 2, 2014

Definitions of CHIMERA

- a single organism composed of genetically distinct cells
  - Proposed methodology is a hybrid of the AIRS approach with Optimal Estimation
- something that exists only in the imagination and is not possible in reality
  - features of the algorithm we have proposed has been an unattainable vision for many, many years
- a monstrous fire-breathing hybrid creature composed of the parts of more than one animal
  - Hopefully not

What are some guiding principles of a climate product for retrievals?

- Requires reprocessing of full dataset
  - By extension, this implies (to me) that AIRS/AMSU and CrIS/ATMS have same spectroscopy and retrieval method
    - Alternatively we could make the radiances look the same
  - IASI/AMSU/MHS in the future (next call?)
  - Incorporate MODIS, AVHRR, and VIIRS in the future
- Community accepted error estimates and/or product characterization (a.k.a., averaging kernels)
  - Requires formal error covariance of the *a-priori*
  - … and formal error covariance of the final products.
- A well characterized *a-priori* suitable for multi-instrument time series

What we plan to do

- Retain the good components of the AIRS science team algorithm (see my “Thoughts on Version 7” presentation on Nov. 15, 2012 – also copied at end of this presentation)
  - Sequential solution using subsets of selected channels
  - Vertical basis functions (not necessarily trapezoids)
  - Geophysical co-variance as part of obs-cal error covariance matrix
  - Use of all sounding assets (microwave and IR imagery in future).
  - Cloud clearing
- Modify the code to formally propagate the error
  - Modify each retrieval step to have a formal a-priori state and associated covariance
  - Propagate the error covariance from step to step
  - Output error covariance (T, q) or averaging kernel (trace gases)
    - Note: outputting one of the two allows derivation of the other
- Allows graceful degradation with decreased information content
  - Avoid ad-hoc adjustments and “switches” in product character

Use the NOAA-Unique CrIS/ATMS Processing System (NUCAPS) code as a starting point

- AIRS/AMSU, IASI/AMSU/MHS, and CrIS/ATMS are processed with literally the same code.
  - Extremely fast compared to other approaches (1 CPU for CrIS/ATMS)
  - Same underlying spectroscopy (as best as we could do)
  - Instrument agnostic: specific items are file-driven, not hardwire
  - Code is backward and forward (as much as possible) compatible.
  - Retrieval components are programmable via namelists (can quickly compare retrieval enhancements and/or methodologies).
  - Operational code is a “filtered” version of the science code.
- **Can use it perform interesting experiments**
  - AIRS + ATMS
  - Jennifer Wei’s O3 tropo-pause relative climatology/retrieval
  - Etc. etc.
- **Can form the basis for other team contributions**

Use the NOAA-Unique CrIS/ATMS Processing System (NUCAPS) code as a starting point

- Note that NUCAPS, by design, is the AIRS Science Team Algorithm
  - It is not my algorithm, it is only my code
  - It is based on the original AIRS science team code
    - Originated from Joel Susskind's science code delivery to JPL (circa 1998) and maintained changes from GSFC (through v5)
    - Worked with Larrabee Strow to obtain SARTA for IASI and CrIS and integrate SARTA into the retrieval framework.
    - Worked with Phil Rosenkranz to integrate and maintain the microwave components
      - Very excited to work with Bjorn's team to continue this work and use his results
    - Worked with Mitch Goldberg to integrate and maintain the regression components
      - In the operational environment the neural net approach was demanding of operational resources, difficult to simultaneously maintain operationally for Aqua, Metop-A, Metop-B, and NPP/JPSS, and worried about it being over-trained w.r.t. global eigenvector regression

So, what a-priori should we use?

- Choice of a-priori is critical
  - Want a prior that contributes unique information in low information content domains.
  - Want a prior that the *climate* community considers to be well behaved

Prior information for T/q	Pro/Con
Climatology	Simple and constant, see retrieval skill
Microwave only O-E w/ clim	With ATMS has high IC, but not Aqua
NCEP Reanalysis	R1 (w/o satellite), f(time) for others
ERA-Interim	System changes with time, good q(p)
MERRA-2	Consistent reprocessing

- Prior for trace gases will be simple climatologies
  - Except ozone, where we will use tropopause relative climatology (Wei 2010 JAOT v.27 p.1123)

Summary of top level difference of some operational systems with CHIMERA

system	Methodology	First guess	*A-priori*	Level-3
AIRS version.5	AST	Regression	None	Yes
AIRS version.6	AST	Neural net	None	Yes
NOAA IASI	AST	Regression	None	No
NUCAPS	AST	Regression	None	No
CRiMSS EDR	O-E	ATMS ret	Climatology	No
CHIMERA	O-E	Regression	Reanalysis	Yes

Validation

- "extraordinary claims require extraordinary evidence" Carl Sagan
  - We should avoid making algorithm choices using the same data sets used in “training” of algorithm or QC components.
  - We should partition improvements into those from null-space and those from physical measurement concepts
- Should the goal be to use IR everywhere?
  - Cloud clearing is known to fail in regions of high moisture or surface variability.
    - Error estimates, including off-diagonal components, will be significantly improved with error propagation through the retrieval steps and use of stable *a-priori*
  - O-E allows for a graceful transition from infrared-dominated to microwave-dominated products in information limited domains.
    - Error propagation and averaging kernels explicitly describes the information content of the sounding information w.r.t. the prior information content.

Validation using in-situ

- We will inter-compare CHIMERA with NUCAPS and AIRS v.6 in order to verify the characteristics of the product
  - These datasets ensure independence of any training (*e.g.*, bias corrections, a-priori)
  - Inter-compare with other methods (e.g., Susskind, Moncet, Irion)

Dataset	Location	Timeframe	Ret. Variables
DOE ARM CART	TWP, SGP, NSA, mobile	~500 RS-92/y, starting 2012	SST, T, WV
AEROSE	Tropical mid-lat Atlantic	~100 RS92 and ~25 O3	SST, T, WV, O3
COSMIC GPS	Global	Yearly	UT T, tropopause
SHADOZ	Southern Hemis.	1998 – present	O3
MOSAIC/IAGOS	Northern Hemis.	1996 – present	WV, CO, O3
HIPPO	Pacific Ocean	2009-2011	CO, CH4, CO2, N2O

Validation with respect to short-term variability indices (year.2)

- Verify algorithm performance with respect to existing community reference networks
  - NCDC ERSST/ICOADS (*)
  - NOAA CPC (+)

Variability Mode	Observed Variables	AIRS/CrIS overlapping domain
El Nino Southern Oscillation (ENSO)	SST (*), Total Cloudiness (*), T(*), q(*), OLR(+)	2012-present Tropical Ocean
Pacific Decadal Oscillation (PDO)	SST (*), Total Cloudiness (*), T(*), q(*), OLR(+)	2012-present Global
Atlantic Multi-decadal Oscillation (AMO)	SST (*), Total Cloudiness (*), T(*), q(*), OLR(+)	2012-present Global

Evaluate with short-term climate sensitivity indices (year.3)

- Assess the ability of this algorithm to measure short-term climate feedbacks

$$\Delta R = dR/dT_s \, \Delta T_s = [\partial R/\partial T_s + \partial R/\partial q \, \partial q/\partial T_s + \partial R/\partial \Gamma \, \partial \Gamma/\partial T_s + \ldots]\Delta T_s +$$

- Use forward model to compute radiative kernels, (dR/dX terms)
- Use CHIMERA products to compute the feedback terms, (dX/dTs terms)
- Verify that X=q and X=OLR over produces physically correct over diverse geophysical regimes for both Aqua and NPP

CHIMERA satisfied the Suomi-NPP Science Team's recommendations for a science-quality algorithm

- Capable of processing CrIS full-resolution spectra (Gambacorta 2013 IEEE GRSL);
- Will produce a satellite climatology for temperature, water vapor, and trace gases from Aqua/AIRS/AMSU and Suomi-NPP CrIS/ATMS
  - also capable of being extended to IASI/AMSU/MHS
  - however, that extension was not being proposed here
- Retrieval approach that elucidates climate signals without bias;
- Designed, from the beginning, to be product-centric rather than sensor-centric;
- Outputs the full-geophysical state and that output can be used to compute radiances;
  - Products include surface, cloud, O_3, CO, CH_4, CO_2, SO_2, HNO_3, and N_2O, products in addition to cloud cleared radiances, temperature and moisture;
- Uses an open framework.
  - other researchers can link other algorithms for the core products and new algorithms for ancillary products (e.g., cloud microphysical products, trace gases, etc.).
- Could add new products
  - Ammonia (NH_3), Formic Acid (HCOOH), and Peroxyacetyl Nitrate (PAN)

QUESTIONS?

Constraints and Assumptions for the AIRS Science Team (AST) Algorithm

- Must be able to process, end-to-end (using ≤ 10 250 MHz CPU's in 2002)
  - NUCAPS does ~1 retrieval per 0.12 seconds on modern CPUs
  - AIRS, IASI, and CrIS all acquire 1 FOR in ~0.27 seconds
- Only static data files can be used
  - One exception: model surface pressure.
  - Cannot use output from model or other instrument data.
  - Maximize information coming from AIRS radiances.
- Cloud clearing will be used to "correct" for cloud contamination in the radiances.
  - Amplification of Noise, A, is a function of scene $0.33 \le A < \approx 5$
  - Spectral Correlation of Noise is a function of scene
  - IR retrievals must be available for all Earth conditions within the assumptions/limitations of cloud clearing.

Flow Diagram of NUCAPS Retrieval Steps

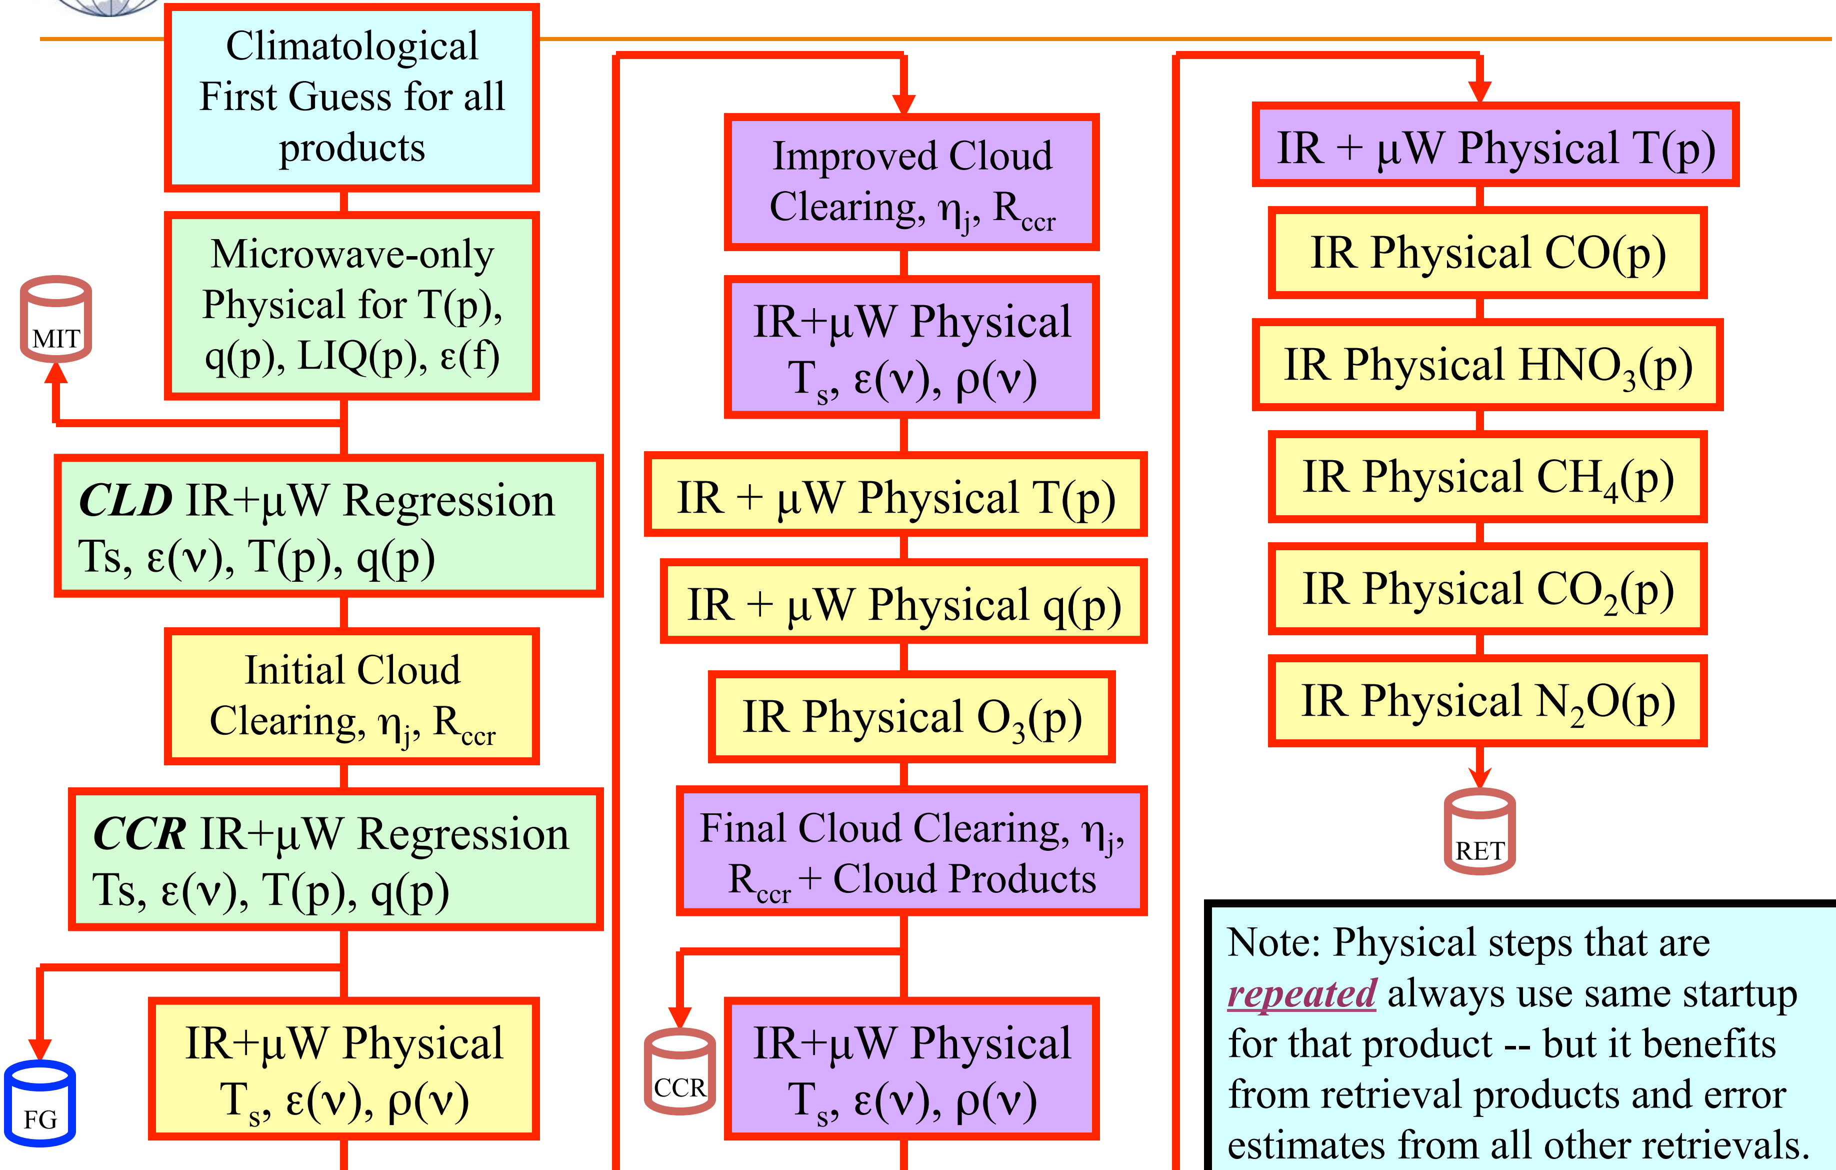

Note: Physical steps that are ***repeated*** always use same startup for that product -- but it benefits from retrieval products and error estimates from all other retrievals.

Summary of products from AIRS, IASI and NUCAPS Algorithm

gas	Range (cm^{-1})	Precision	d.o.f.	Interfering Gases	Sensitivity
T	650-800 2375-2395	1.5K/km	6-10	H2O,O3,N2O emissivity	surface to ~1 mb
H_2O	1200-1600	15%	4-6	CH4, HNO3	surf to 300 mb
Cloud P, T, fraction	700-900	25 mbar, 1.5K, 5%	≈2	CO2, H2O	surface to tropopause
O_3	1025-1050	10%	1+	H2O,emissivity	Lower strat.
CO	2080-2200	15%	≈ 1	H2O,N2O	Mid-trop
CH_4	1250-1370	1.5%	≈ 1	H2O,HNO3,N2O	Mid-trop
CO_2	680-795 2375-2395	0.5%	≈ 1	H2O,O3 T(p)	Mid-trop
Volcanic SO_2	1340-1380	50% ??	< 1	H2O,HNO3	flag
HNO_3	860-920 1320-1330	50% ??	< 1	emissivity H2O,CH4,N2O	Upper trop
N_2O	1250-1315 2180-2250	5% ??	< 1	H2O H2O,CO	Mid-trop

Thoughts on Version 7

NASA Sounder Team Meeting

(NOTE: This presentation draws on some conclusions shown in the previous presentation (CrIMSS EDR status)

Christopher Barnet
NOAA/NESDIS/STAR

Nov. 15, 2012

Objective

- This is a philosophical presentation intended to incite discussion on potential version 7 systems.
    - Primary concern is that users may not be aware of subtle characteristics of our products.
    - Primary goals are (1)encourage community acceptance of the AIRS products and (2) further exploit AIRS information content.
    - It is also possible that v7 could contain multiple product types (*e.g.*, one for climate and one for weather applications.
  - My opinion of product attributes is not intended to offend any algorithm developer
    - Although, maybe it is more accurate to say I am trying to offend all algorithm developers equally.
- This discussion is at a high level (*i.e.*, no equations)
  - But, obviously, a primary objective of this talk is to discuss options in a mathematically rigorous manner.

1DVAR versus AIRS Science Team Method

Simultaneous (1DVAR)	Sequential (AIRS method)
Solve all parameters simultaneously	Solve each state variable (*e.g.*, T(p)), separately.
Error covariance includes only instrument model.	Error covariance is computed for all *relevant* state variables that are held fixed in a given step. Retrieval error covariance is propagated between steps.
Each parameter is derived from all channels used (*e.g.*, can derive T(p) from CO2, H2O, O3, CO, … lines).	Each parameter is derived from the best channels for that parameter (*e.g.,* derive T(p) from CO2 lines, q(p) from H2O lines, etc.)
A-priori must be rather close to solution, since state variable interactions can de-stabilize the solution.	*A-priori* can be less complex for sequential with well selected channels.
Regularization must include *a-priori* statistics to allow mathematics to separate the variables and stabilize the solution.	Regularization can be reduced (smoothing terms) and does not require *a-priori* statistics for most geophysical regimes.
This method has large state matrices (all parameters) and covariance matrices (all channels used). Inversion of these large matrices is computationally expensive.	State matrices are small (largest is 25 T(p) parameters) and covariance matrices of the channels subsets are quite small. Very fast algorithm. Encourages using more channels.
Has never been done simultaneously with clouds, emissivity(ν), SW reflectivity, surface T, T(p), q(p), O3(p), CO(p), CH4(p), CO2(p), HNO3(p), N2O(p)	*In-situ* validation and satellite inter-comparisons indicate that this method is robust and stable.

Simplified Flow Diagram of the AIRS Science Team Algorithm

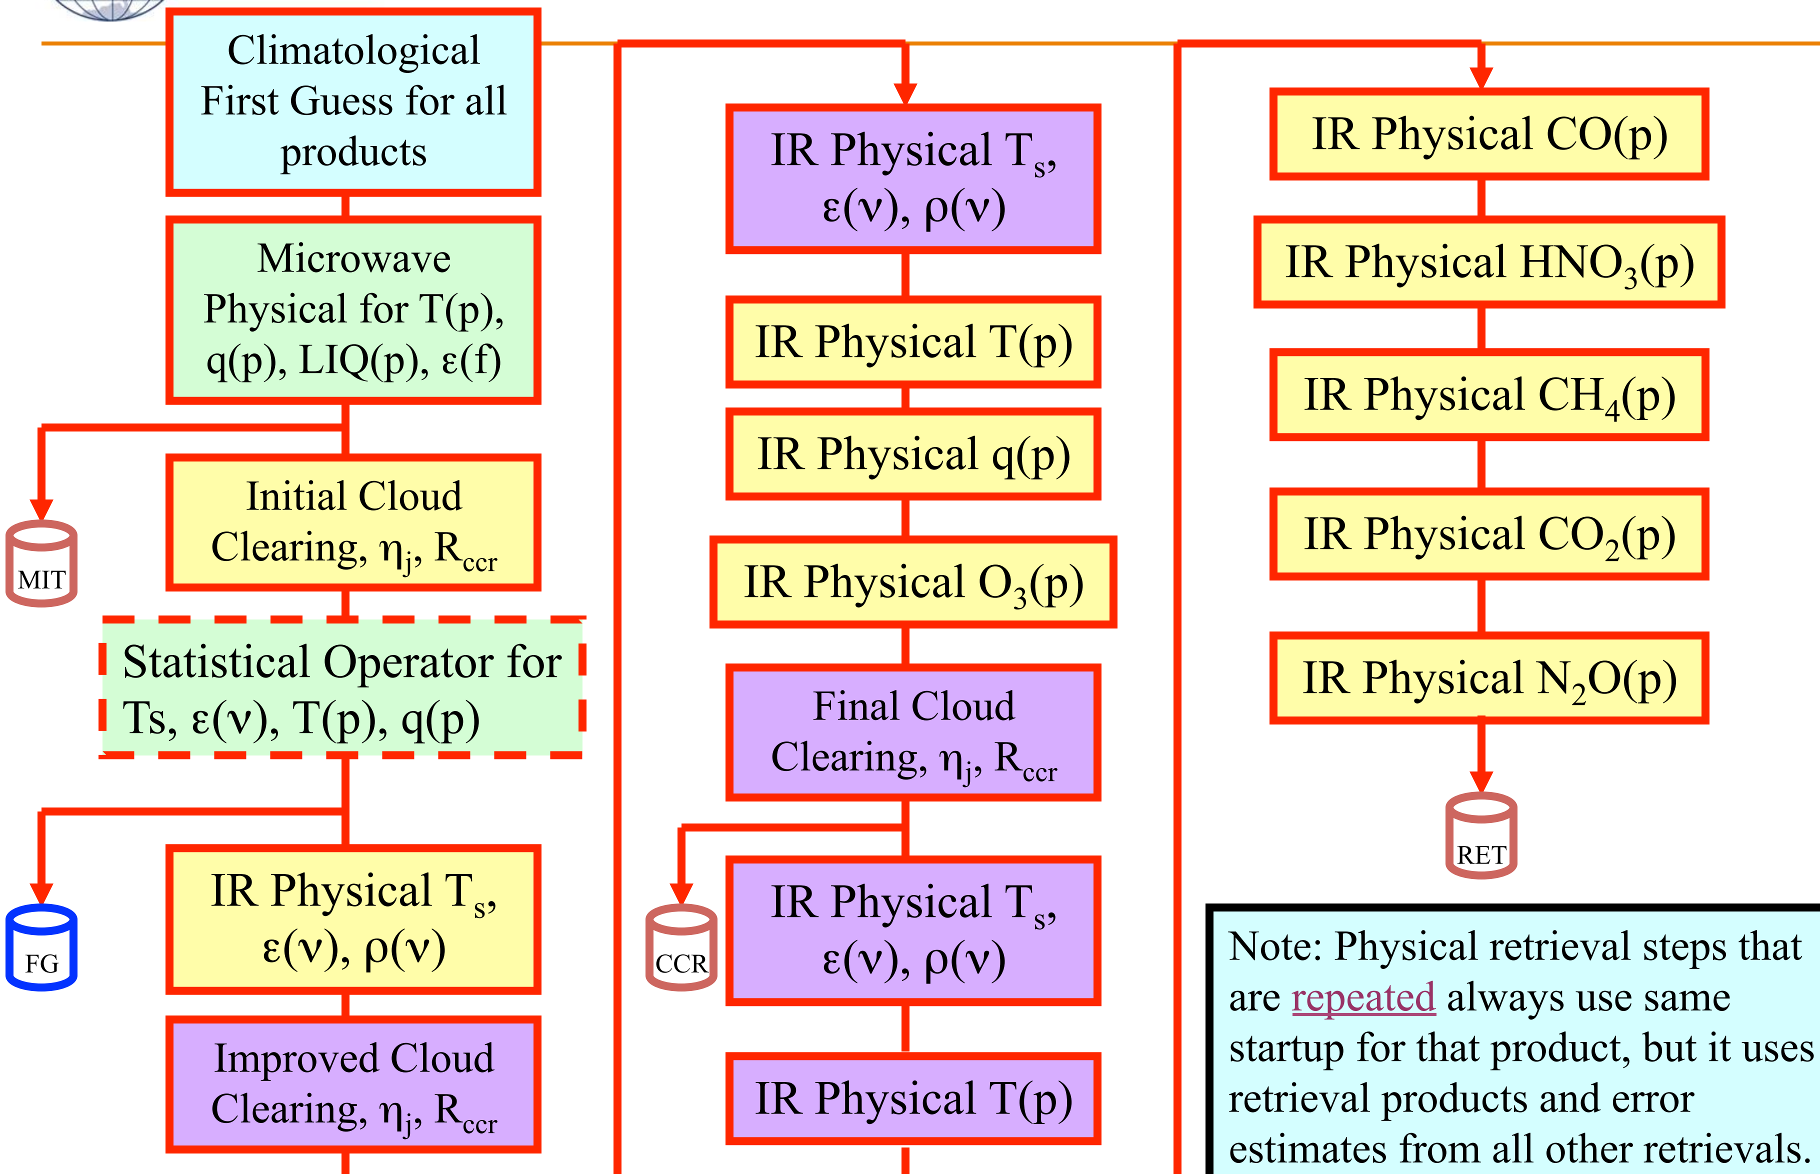

Note: Physical retrieval steps that are repeated always use same startup for that product, but it uses retrieval products and error estimates from all other retrievals.

Advantages of the AIRS Approach

- Sequential physical algorithm allows for a robust and stable system with minimal prior information
  - Sequential approach allows the more linear parameters to be solved for first -- can make the algorithm very stable
  - Can solve for all significant signals in the AIRS radiances.
- Error from previous steps are mapped into an error estimate from interfering parameters
  - A unique feature of this algorithm is that error estimates from previous steps are mapped into subsequent steps
  - The observation covariance (S_ε in Rodgers 2000) contains both on- and off-diagonal terms composed of $(dR/dX)\cdot\delta x$ for all x's that are considered interference (including cloud clearing, correlation due to apodization, etc.).
  - Can be more robust than simultaneous retrieval because each step uses optimal sampling of channels (*i.e.*, low interference).

Advantages of optimal estimation

- O-E explicitly constrains the answer to lie within expectation of reasonable answers
  - Prior assumptions are always implicit in any retrieval approach
  - Note that "reasonable" can be in the eye of the beholder and sometimes that means a preference in the vertical null space.
- O-E explicitly derives the answer from prior information
  - in this sense, 1st guess can only speed up convergence
  - with enough iterations the same answer is usually achieved (up to non-linearity of Jacobians)
- Information content (or errors) in retrieval state can be partitioned between instrument and prior contributions
  - Averaging kernels or error covariance have more value

Graphical representation of O-E

Contour of the prior PDF for a 3-D state, S_a

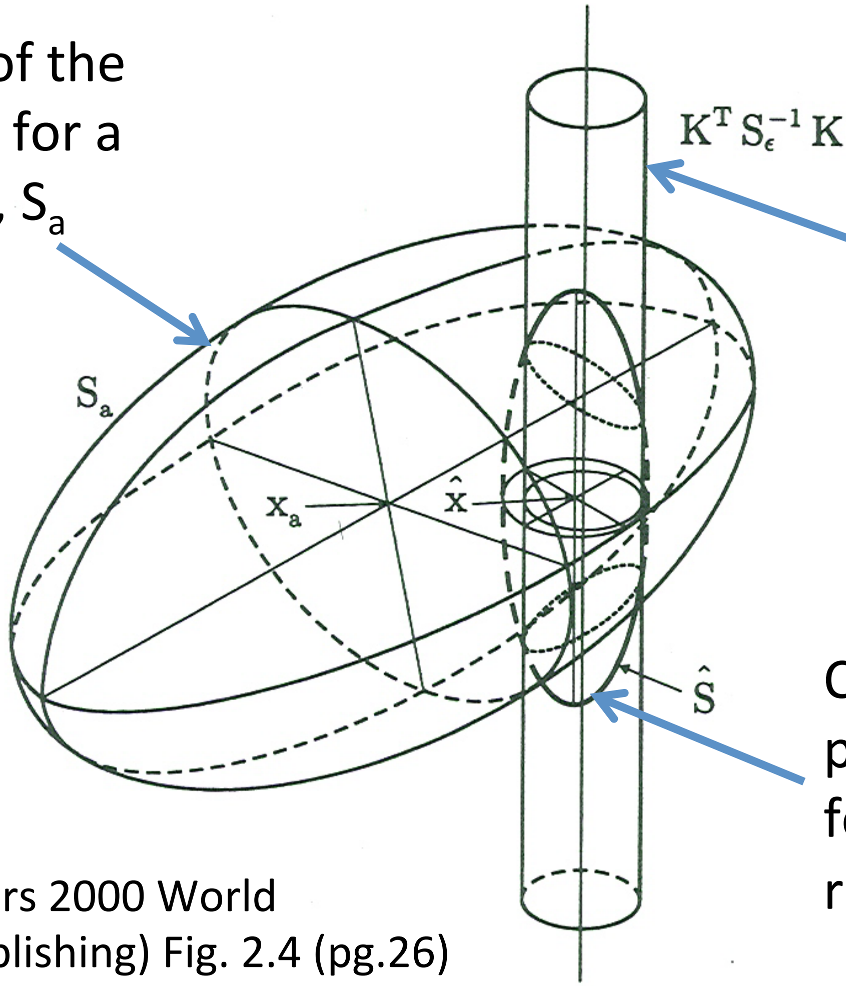

2-D measurement (i.e., no sensitivity to 3rd dimension) mapped into state space

Contour of the posterior PDF for a optimal retrieval.

From (Rodgers 2000 World Scientific Publishing) Fig. 2.4 (pg.26)

Statistical Operators

- Statistical retrievals are those that *fit* radiances, R, directly to an ensemble of geophysical parameters, X
  - X = f(R), usually all radiances are used
  - Neural net: $X = A*\alpha(R) + B*\beta(R) + C*\gamma(R) + D*\delta(R)$
  - Linear regression: $X = A*R$
  - Neural Net has more free degrees of freedom
- Information can be derived from correlations
  - e.g. when we used to have an ozone regression we found that tropospheric ozone was being derived from AIRS channels sensitive to tropopause height and carbon monoxide
    - Would we call this a "measurement" or is it an "index“
    - We did learn from this – led to tropopause relative first guess

Training of AIRS statistical operators (global versus regional)

- NOAA regression was trained globally and used eigenvector regularization
  - We wanted to constrain the degrees of freedom allowed
  - 80 PCs with stratification into 4 view angle bins
- Neural net trained regionally, 200+ stratifications
  - 2 ascending and descending
  - 3 latitude bands (N.H., temperate, S.H.)
    - Each has frozen/non-frozen ocean, 5-7 surface pressure over land
  - 4 seasons
  - Version 6 Neural Net has *significantly* more free degrees of freedom to “fit” ECMWF
- Therefore, the differences between NOAA linear regression to MIT neural network approach can be do with these choices in stratification, constraints, etc.

Training of Statistical Operators (Geophysical Variance)

- Training must include every condition seen on Earth over the lifetime of the mission
  - For example, early in the AIRS mission we had issues with volcanic SO2 from Etna
    - volcanic SO2 was not in our early training (now it is)
    - Statistical operator extrapolated to completely unrealistic profiles
    - When it is good, it is very very good, but when it is bad ....
- Sub-resolved structure, being derived by correlations, needs expansive training
  - using ECMWF for training means we build in all ECMWF errors of the day
    - e.g., ECMWF ozone in May 2012 has very large errors
    - if this had been used for training of an ozone product it would have caused erroneous ozone products
- I would argue that there can never be enough training
  - Are there less obvious attributes of ECMWF that we have inadvertently embedded into our product?

Some concerns with the statistical operator have already been raised

- Vertical structure has been shown to be greater than that which we can measure (Larrabee, Oct. 2011 AIRS meeting)
  - Statistical operator has ability to relate sub-resolved structure with AIRS radiances.
  - When the wrong structure is imposed in our first guess it is not removed by our physical retrieval (to be discussed in a few slides)
- Eric Maddy has shown that while Version 6 has significantly better statistics for temperature and water vapor profiles the cloud cleared radiance statistics are identical to v5.9
  - Implies that the improvement in T(p) RMS may be due to sub-resolved vertical structures (*i.e.*, improvements in our null space, not our measurement)

Some mathematical issues with AIRS physical retrieval methodology

- We do not have a formal a-priori constraint.
  - We do have an ad-hoc "background term"
    - back in the day, I had convinced myself it provided the same functionality as a Roger's background term (recursively)
    - but this is not true, it does not equate to minimization of a cost function
  - iterations are done w.r.t. previous state, with some % held back
    - advantage: this retains the full vertical structure of first guess
    - disadvantage: there is no constraint, physical retrieval believes first guess
      - even if we characterized the statistical operator's covariance that information would not be used by our physical retrieval
- We only map the diagonal component of the error covariance into down-stream steps.
  - Eric Maddy has shown there is a robust way to pass the full covariance from one step to the next (Mar. 23 2007 AIRS meeting, my talk in session 6 and Eric Maddy Apr. 27, 2011 session 6)
- The physical algorithm has become a "QC" of the statistical operator
  - The goal is to select as many "good" cases and reject the "bad" regions
  - Usually, the statistical operator is very good (better than we can measure) so that "best" physical retrieval is one that does nothing
    - Tendency to over-regularize the physical retrieval

So, what is the most desirable system?

- If we fixed the "background term" then we must select a real prior state (need both state and covariance)
  - This can be non-trivial: for some products (or simultaneous "1DVAR-like" covariance) the covariance could be very difficult to construct.
  - Note that model priors also contain information on dynamics

Prior information	Potential User Community
Statistical (with covariance)	Regional NWP
Climatology	Process studies
Forecast Model X (w/o AIRS R)	Global NWP for X, X=GFS,ECM,GMAO,etc.
Re-analysis product X (w/o AIRS R)	Historical climate for X

- O-E can also be done sequentially (and with cloud clearing) but for meaningful error estimates (or Averaging Kernels) we will need to improve the propagation of the error covariance downstream
- And there is a choice between clear-FOV retrievals (low daily yield, very good error characterization) or cloud clearing (high yield, complex error characteristics).

We could add more information content
(*i.e.*, minimize dependence on prior information)

- MODIS radiances
  - NOAA already has MODIS IR convolved to AIRS FOVs
    - We also have AVHRR IR convolved to IASI (to be installed 2012)
    - … and will have VIIRS IR convolved to CrIS (to be installed 2014)
  - It improves cloud clearing (part of our phase-2 IASI system)
  - Could potentially improve surface retrieval
- With degradation of AMSU and loss of HSB consider using alternative microwave radiances
  - CrIS/ATMS results demonstrate that the microwave information is important, especially for moisture
  - Eric has run ATMS+AIRS
    - Quick look results imply that the increase of information content may be more valuable than degradation of co-location
  - We could employ NOAA AMSU over life of AIRS mission

Validation

- "extraordinary claims require extraordinary evidence" Carl Sagan
  - We should avoid making algorithm choices using the same data sets used in “training” of algorithm or QC components.
  - We should partition improvements into those from null-space and those from physical measurement concepts
- Should the goal be to use IR everywhere?
  - Cloud clearing is known to fail in regions of high moisture or surface variability and has large non-Gaussian errors when it fails.
  - There is a trade-off between quality and robustness as scenes becomes more complicated.
  - CrIMSS metric is to have a retrieval everywhere
    - We look at both MW-only and IR+MW rets and decide where the IR retrievals have better performance.
    - To do this we must look at both accepted and rejected IR retrievals
    - We also require validation of a full profile (from TOA to surface).

Backup: O-E vs AIRS equations (somewhat simplified to make them look similar)

O-E pivoting off of prior state:

$$X_j^i = X_j^A + \left[K_{j,n}^T \cdot N_{n,n}^{-1} \cdot K_{n,j} + C_{j,j}^{-1}\right]^{-1} \cdot K_{j,n}^T \cdot N_{n,n}^{-1} \cdot \left[R_n^{obs} - R_n(X^{i-1}) + K_{n,j} \cdot \left(X_j^{i-1} - X_j^A\right)\right]$$

Minimizes the cost function:

$$J = \left(R_n^{obs} - R_n\left(X_j^{i-1}\right)\right)^T \cdot N_{n,n}^{-1} \cdot \left(R_n^{obs} - R_n\left(X_j^{i-1}\right)\right) + \left(X_j^{i-1} - X_j^A\right)^T \cdot C_{j,j}^{-1} \cdot \left(X_j^{i-1} - X_j^A\right)$$

Equivalent to pivoting off of the previous iteration:

$$X_j^i = X_j^{i-1} + \left[K_{j,n}^T \cdot N_{n,n}^{-1} \cdot K_{n,j} + C_{j,j}^{-1}\right]^{-1} \cdot \left[K_{j,n}^T \cdot N_{n,n}^{-1} \cdot \left(R_n^{obs} - R_n(X^{i-1})\right) - C_{j,j}^{-1} \cdot \left(X_j^{i-1} - X_j^A\right)\right]$$

AIRS Science Team approach:

$$X_j^i = X_j^{i-1} + \left[K_{j,n}^T \cdot N_{n,n}^{-1} \cdot K_{n,j} + H_{j,j}\right]^{-1} \cdot K_{j,n}^T \cdot N_{n,n}^{-1} \cdot \left[R_n^{obs} - R_n(X^{i-1}) - \Psi_n^{i-1}\right]$$

H is a smoothing constraint and the background term is derived with respect to unregularized (LSQ) retrieval

$$\Psi_n^{i-1=0} = 0$$

$$\Psi_n^{i-1} = K_{n,j}^{i-1} \cdot \left(X_j^{i-1}(0) - X_j^{i-1}\right)$$

$$X_j^i(0) = X_j^{i-1} + \left[K_{j,n}^T \cdot N_{n,n}^{-1} \cdot K_{n,j}\right]^{-1} \cdot K_{j,n}^T \cdot N_{n,n}^{-1} \cdot \left[R_n^{obs} - R_n(X^{i-1}\right]$$

Discussion

- Suggested Rules for Engagement
  - suspend judgment
  - no speeches (1 minute rule)
  - one person speaks at a time (one idea at a time)
  - no killer phrases
  - hitchhiking is okay
  - be creative

All day long, a tough gang of astrophysicists would monopolize the telescope and intimidate the other researchers.